# Multi-Residual Networks: Improving the Speed and Accuracy of Residual Networks

Masoud Abdi, and Saeid Nahavandi, *Senior Member, IEEE*

arXiv:1609.05672v4 [cs.CV] 15 Mar 2017

***Abstract*—In this article, we take one step toward understanding the learning behavior of deep residual networks, and supporting the observation that deep residual networks behave like ensembles. We propose a new convolutional neural network architecture which builds upon the success of residual networks by explicitly exploiting the interpretation of very deep networks as an ensemble. The proposed multi-residual network increases the number of residual functions in the residual blocks. Our architecture generates models that are wider, rather than deeper, which significantly improves accuracy. We show that our model achieves an error rate of 3.73% and 19.45% on CIFAR-10 and CIFAR-100 respectively, that outperforms almost all of the existing models. We also demonstrate that our model outperforms very deep residual networks by 0.22% (top-1 error) on the full ImageNet 2012 classification dataset. Additionally, inspired by the parallel structure of multi-residual networks, a model parallelism technique has been investigated. The model parallelism method distributes the computation of residual blocks among the processors, yielding up to 15% computational complexity improvement.**

***Index Terms*—Deep residual networks, convolutional neural networks, image classification, deep learning.**

## I. Introduction

CONVOLUTIONAL neural networks [18] have contributed to a series of advances in tackling image recognition and visual understanding problems [17], [25], [37]. They have been applied in many areas of engineering and science [34], [22], [15]. Increasing the network depth is known to improve the model capabilities, which can be seen from AlexNet [17] with 8 layers, VGG [26] with 19 layers, and GoogleNet [32] with 22 layers. However, increasing the depth can be challenging for the learning process because of the vanishing/exploding gradient problem [11], [2]. Deep residual networks [8] avoid this problem by using identity skip-connections, which help the gradient to flow back into many layers without vanishing. The identity skip-connections facilitate training of very deep networks up to thousands of layers that helped residual networks win five major image recognitions tasks in ILSVRC 2015 [24] and Microsoft COCO 2015 [21] competitions.

However, an obvious drawback of residual networks is that every percentage of improvement requires significantly increasing the number of layers, which linearly increases the computational and memory costs [8]. On CIFAR-10 image classification dataset, deep residual networks with 164-layers and 1001-layers reach a test error rate of $5.46\%$ and $4.92\%$ respectively, while the 1001-layer has six times more computational complexity than the 164-layer. On the other hand, wide residual networks [36] have 50 times fewer layers while outperforming the original residual networks. It seems that the power of residual networks is due to the identity skip-connections rather than extremely increasing the network depth.

Nevertheless, a recent study supports that deep residual networks act like ensembles of relatively shallow networks [33]. This is achieved by showing the existence of exponential paths from the output layer to the input layer that gradient information can flow. Also, observations show that removing a layer from a residual network, during the test time, has a modest effect on its performance. Additionally, it shows that most of the gradient updates during optimization come from ensembles of relatively shallow depth. Moreover, residual networks do not resolve the vanishing gradient problem by preserving the gradient through the entire depth of the network. Instead, they avoid the problem by ensembling exponential networks of different length. This raises the importance of *multiplicity* that refers to the number of possible paths from the input layer to the output layer [33].

Inspired by these observations, we introduce multi-residual networks (Multi-ResNet) which increase the multiplicity of the network, while keeping its depth fixed. This is achieved by increasing the number of residual functions in each residual block. We then show that the accuracy of a shallow multi-residual network is similar to a deep 110-layer residual network. This supports that deep residual networks behave like ensembles instead of a single extremely deep network. Next, we examine the importance of *effective range* which is the range of paths that significantly contribute towards gradient updates.

We show that for a residual network deeper than a threshold $n_0$, increasing the number of residual functions leads to a better performance than increasing the network depth. This leads to a lower error rate for the multi-residual network with the same number of convolutional layers as the deeper residual network. Experiments on ImageNet, CIFAR-10, and CIFAR-100 datasets show that multi-residual networks improve the accuracy of deep residual networks and outperform almost all of the existing models.

We demonstrate that a 101-layer Multi-ResNet with two residual functions in each block outperforms the top-1 accuracy rate of a 200-layer ResNet by $0.22\%$ on the ImageNet 2012 classification dataset [24]. Also, using moderate data augmentation (flip/translation), multi-residual networks achieve an error rate of $4.35\%$ and $20.42\%$ on CIFAR-10 and

M. Abdi and S. Nahavandi are with the Institute for Intelligent Systems Research and Innovation (IISRI), Deakin University , Waurn Ponds, VIC 3216, Australia. e-mail: mabdi@deakin.edu.au.

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

CIFAR-100 receptively (based on five runs). This is **6%** and **10%** improvement compared to the residual networks with identity mappings [10] with almost the same computational and memory complexity. The proposed multi-residual network achieves a test error rate of $3.73\%$ and $19.45\%$ on CIFAR-10 and CIFAR-100.

Concurrent to our work, ResNeXt [35] and PolyNet [38] achieved second and third place in the ILSVRC 2016 classification task[1]. Both models increase the number of residual functions in the residual blocks similar to our model, while PolyNet inserts higher order paths into the network as well.

Eventually, a model parallelism technique has been explored to speed up the proposed multi-residual network. The model parallelism approach splits the calculation of each block between two GPUs, thus each GPU can simultaneously compute a portion of residual functions. This leads to the parallelization of the block, and consequently the network. The resulting network has been compared to a deeper residual network with the same number of convolutional layers that exploits data parallelism. Experimental results show that in addition to being more accurate, multi-residual networks can also be up to **15%** faster.

In summary, the contributions of this research are:

- We take one step toward understanding deep residual networks and supporting that deep residual networks behave like ensembles of shallow networks, rather than a very deep network.
- Through a series of experiments, we show the importance of the effective range in residual networks, which is the range of ensembles that significantly contribute toward gradient updates during optimization.
- We introduce multi-residual networks that is shown to improve the classification accuracy of deep residual networks and many other state-of-the-art models.
- We propose a model parallelism technique that is able to reduce the computational complexity of the multi-residual networks.

The rest of the paper is organized as follows. Section II details deep residual networks and other models capable of improving the original residual networks. The hypothesis that residual networks are exponential ensembles of relatively shallow networks is explained in Section III. The proposed multi-residual networks and the importance of the effective range are discussed in Section IV. Supporting experimental results are presented in Section V. Concluding remarks are provided in Section VI. A pre-print version of this paper [1] is available at https://arxiv.org/abs/1609.05672, and the code to reproduce the results can be found at https://github.com/masoudabd/multi-resnet.

## II. Related Work

A residual block consists of a residual function $f$, and an identity skip-connection (see Figure 2), where $f$ contains convolution, activation (ReLU) and batch normalization [14] layers in a specific order. In the most recent residual network the order is normalization-ReLU-convolution which is known as the pre-activation model [10].

Deep residual networks contain many stacked residual blocks with $y = x + f(x)$, where $x$ and $y$ are the input and output of the block. Moreover, a deep residual network with the identity skip-connections [10] can be represented as:

$$x_{l+1} = x_l + f_{l+1}(x_l) \tag{1}$$

where $x_l$ is the input of $l^{\text{th}}$ residual block, and $f_l$ contains the weight layers. Additionally, Highway Networks [31], [30] also employ parametrized skip-connections that are referred to as *information highways*. The skip-connection parameters are learned during training, which control the amount of information that can pass through the skip-connections.

Residual networks with stochastic depth [13] use Bernoulli random variables to randomly disable the residual blocks during the training phase. This results in a shallower network at the training phase, while having a deeper network at the test phase. Deep residual networks with stochastic depth improve the accuracy of deep residual networks with constant depth. This is because of the reduction in the network depth which strengthens the back-propagated gradients of the earlier layers, and because of ensembling networks of different depths.

Swapout [27] generalizes dropout [29] and networks with stochastic depth [13] using $px+qF(x)$, where $p$ and $q$ are two Bernoulli random variables. Swapout has the ability to sample from four network architectures $\{0, x, F(x), x + F(x)\}$, therefore having a larger domain for ensembles. Wide residual networks [36] increase the number of convolutional filters, and are able to yield a better performance than the original residual networks. This suggests that the power of residual networks originate in the residual connections, as opposed to extremely increasing the network depth. DenseNet [12] uses a dense connection pattern among the convolutional layers, where each layer is directly connected to all preceding layers.

## III. Deep Residual Networks behave like Ensembles

Deep residual networks [8] are assumed to resolve the problem of vanishing gradients using identity skip-connections that facilitate training of deep networks up to 1202 layers. Nonetheless, recent studies support that deep residual networks do not resolve the vanishing gradient problem by preserving the gradient flow through the entire depth of the network. Instead, they avoid the problem simply by ensembling exponential networks together [33].

Consider a residual network with three residual blocks, and let $x_0$ and $x_3$ be the input and output respectively, applying Equation 1 iteratively gives:

$$\begin{aligned} x_3 &= x_2 + f_3(x_2) \\ &= \Big[x_1 + f_2(x_1)\Big] + f_3(x_1 + f_2(x_1)) \\ &= \Big[x_0 + f_1(x_0) + f_2(x_0 + f_1(x_0))\Big] \\ &+ f_3(x_0 + f_1(x_0) + f_2(x_0 + f_1(x_0))) \end{aligned} \tag{2}$$

A graphical view of Equation 2 is presented in Figure 1a. It is clear that data flow along the exponential paths from the

[1] http://image-net.org/challenges/LSVRC/2016/results

input to the output layer. In other words, every path is a unique configuration that either computes a particular function $f_l(l = 1, \ldots, n)$ or skips it. Therefore, the total number of possible paths from the input to the output is $2^n$, where $n$ is the number of residual blocks. This term is referred to as the *multiplicity* of the network. Furthermore, a residual network can be viewed as a very large implicit ensemble of many networks with different length.

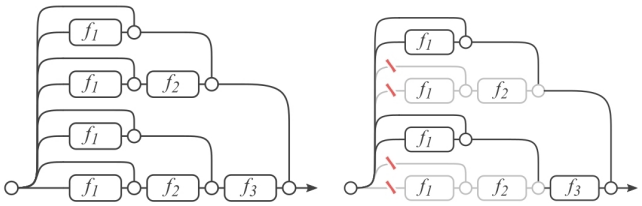


Fig. 1: (a) A residual network; (b) Deleting $f_2$ from a residual network [33]. It can be seen that residual networks have $2^n$ paths connecting the input to the output. Deleting a block from the residual network reduces the number of paths to $2^{n-1}$.

Deep residual networks are resilient to dropping and reordering the residual blocks during the test phase. More precisely, removing a single block from a 110-layer residual network, during the test phase, has a negligible effect on its performance. Whereas, removing a layer from the traditional network architectures, such as AlexNex[17] or VGGnet[26], dramatically hurts the performance of the models (test error more than $80\%$) [33]. This supports the existence of exponential paths from the input to the output layer. Moreover, removing a single residual block during the test phase reduces the number of paths from $2^n$ to $2^{n-1}$ (see Figure 1b),

Additionally, shallow ensembles contribute significantly to the gradient updates during optimization . In other words, in a 110-layer residual network, most of the gradient updates come from paths with only 10-34 layers, and deeper paths do not have significant contribution towards the gradient updates. These are called the *effective paths*, which are relatively shallow compared to the network depth [33].

In order to verify the claim pertaining to the shallow ensembles, one can see that individual paths in a deep residual network have a binomial distribution, where the number of paths with length $k$ is $\binom{n}{k} = \frac{n!}{k!(n-k)!}$. On the other hand, it has been known that the gradient magnitude, during back-propagation, decreases exponentially with the number of functions it goes through [11], [2]. Therefore, the total gradient magnitude contributed by paths of each length can be calculated by multiplying the number of paths with that length, and the expected gradient magnitude of the paths with the same length [33].

Accordingly, a residual network trained with only effective paths has a comparable performance with the full residual network [33]. This is achieved by randomly sampling a subset of residual blocks for each mini-batch, and forcing the computation to flow through the selected blocks only. In this case the network can only see the effective paths that are relatively shallow, and no long path is used.

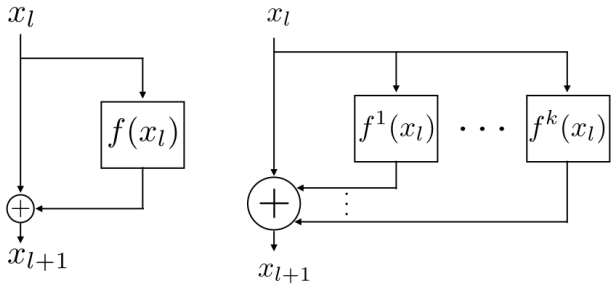


Fig. 2: A residual block (left) versus a multi-residual block (right).

## IV. Multi-Residual Networks

Based on the aforementioned observations, we propose *multi-residual networks* that aim to increase the multiplicity of the residual network, while keeping the depth fixed. The multi-residual network employs multiple residual functions, $f^i$, instead of one function for each residual block (see Figure 2). As such, a deep multi-residual network with $k$ functions has:

$$x_{l+1} = x_l + f^1_{l+1}(x_l) + f^2_{l+1}(x_l) + \cdots + f^k_{l+1}(x_l) \tag{3}$$

where $f^i_l$ is the $i^{th}$ function of the $l^{th}$ residual block. Expanding Equation 3 for $k = 2$ functions and three multi-residual blocks gives:

$$\begin{aligned}
x_3 =& x_2 + f^1_3(x_2) + f^2_3(x_2) \\
=& \Big[x_1 + f^1_2(x_1) + f^2_2(x_1)\Big] + \Big[f^1_3(x_1 + f^1_2(x_1) + f^2_2(x_1))\Big] \\
&+ \Big[f^2_3(x_1 + f^1_2(x_1) + f^2_2(x_1))\Big] \\
=& \Big[x_0 + f^1_1(x_0) + f^2_1(x_0) + f^1_2(x_0 + f^1_1(x_0) + f^2_1(x_0)) \\
&+ f^2_2(x_0 + f^1_1(x_0) + f^2_1(x_0))\Big] \\
&+ \Big[f^1_3(x_0 + f^1_1(x_0) + f^2_1(x_0) + f^1_2(x_0 + f^1_1(x_0) + f^2_1(x_0)) \\
&+ f^2_2(x_0 + f^1_1(x_0) + f^2_1(x_0)))\Big] \\
&+ \Big[f^2_3(x_0 + f^1_1(x_0) + f^2_1(x_0) + f^1_2(x_0 + f^1_1(x_0) + f^2_1(x_0)) \\
&+ f^2_2(x_0 + f^1_1(x_0) + f^2_1(x_0)))\Big]
\end{aligned} \tag{4}$$

It can be seen that the number of terms in Equation 4 is exponentially more than the number of terms in Equation 2. Specifically, in a multi-residual block with $k = 2$ residual functions, the gradient flow has four possible paths: (1) skipping both $f^1$ and $f^2$, (2) skipping $f^1$ and performing $f^2$, (3) skipping $f^2$ and performing $f^1$, (4) performing both $f^1$ and $f^2$. Therefore, the multiplicity of the multi-residual network with two residual functions is $4^n$. In other words, the multiplicity of a multi-residual network with $k$ residual functions and $n$ multi-residual blocks is $2^{kn}$. This is because every function can be either computed or otherwise, giving a

multiplicity of $2^k$ for a block, and a total multiplicity of $2^{kn}$ for the multi-residual network.

### A. Residual Networks Behave Like Ensembles

Based on multi-residual networks, we show that residual networks behave like ensembles. A shallow multi-residual network with the same number of parameters as a 110-layer residual network is able to achieve the accuracy of the residual network. This supports the hypothesis that residual networks behave like exponential ensembles of shallow networks, rather than a single deep network.

| method | depth | k | #params | CIFAR-10(%) |
|---|---|---|---|---|
| resnet[8] | 110 | 1 | 1.7M | 6.61 |
| pre-resnet[10] | 110 | 1 | 1.7M | 6.37 |
| multi-resnet | 8 | 23 | 1.7M | 7.37 |
| [ours] | 14 | 10 | 1.7M | 6.42 |

TABLE I: Classification error on CIFAR-10 test set. A shallow multi-residual network is able to approximate the accuracy of a 110-layer residual network.

A Multi-ResNet with the depth of 8 and $k = 23$ residual functions, and a Multi-ResNet with the depth of 14 and $k = 10$ residual functions are trained. Both networks have roughly the same number of parameters, which is the same as those in the 110-layer residual network. The networks are trained with the same hyper-parameters and training policy as in [8]. Table I summarizes the test errors on CIFAR-10. It can be seen that the classification accuracy of the shallow multi-residual network with 14-layer depth almost reaches that of the 110-layer residual network.

### B. The Effective Range

Based on the observation that residual networks behave like ensembles of shallow networks, a question is posed: **what is the relationship between the range of the effective paths and the depth of the residual network?** More precisely, what is the relationship between the effective range of a residual network with $n$ residual blocks and that of a residual network with $cn$ residual blocks, where $c$ is a constant number?

We hypothesize that this relationship is **not linear**. This implies that if the effective range of a residual network with $n$ blocks is $[a, b]$, the effective range of a residual network with $cn$ blocks is not $[ca, cb]$. Instead, it is shifted and/or scaled toward shallower networks. This is because of the exponential reduction in the gradient magnitude [33], [2]. Eventually, the upper bound of the effective range is lower than $cb$. This could be a potential reason for the problem that every percentage of improvement in deep residual networks requires significantly increasing the number of layers.

### C. Residual Networks versus Multi-Residual Networks

Consider a residual network $R$ with $n$ residual blocks, and let $c$ be a constant integer. We would like to construct two residual networks by: (1) increasing the number of residual blocks to $cn$, which results in a residual network with $c$ times depth of $R$ (excluding the first and last layers), (2) retaining the same depth while increasing the number of residual functions by $c$. The number of parameters of the subsequent networks are roughly the same. One can also see that the multiplicity of both networks are $2^{cn}$, but **how about the effective range of (1) and (2)**?

As discussed in the previous part, the effective range of (1) does not increase linearly, whereas the effective range of (2) increases linearly due to the increase in the residual functions. This is owing to the increase in the number of paths of each length, which is a consequence of changing the binomial distribution to a multinomial distribution. Note that this analysis holds true for $n \geq n_0$, where $n_0$ is a threshold; otherwise the power of the network depth is clear both in theory [7], [6], [5] and in practice [17], [26], [32].

## V. Experimental Results

To support our analyses and show the effectiveness of the proposed multi-residual networks, a series of experiments has been conducted on CIFAR-10 and CIFAR-100 datasets. Both datasets contain 50,000 training samples and 10,000 test samples of $32 \times 32$ color images, with 10 (CIFAR-10) and 100 (CIFAR-100) different categories. We have used *"moderate data augmentation"* (flip/translation) as in [10], and training is done using stochastic gradient descent for 200 epochs with a weight decay of $10^{-4}$ and momentum of 0.9 [8]. The network weights have been initialized as in [9].

### A. The Effective Range Phenomena

Consider a pre-activation version of the residual network with the basic-blocks [10]. Three pairs of residual network and multi-residual network are trained. The residual network is $k$ times deeper than the corresponding multi-residual network (excluding the first and last layers). On the other hand, the multi-residual network computes $k$ residual functions. A residual block might be removed to compensate the difference in the number of parameters to form a fair comparison between the pairs. The median of five runs with mean±std in the parentheses are reported in Table II. Test error curves are also depicted in Figure 3, where each curve is the mean of five runs. All networks are trained with the same hyper-parameters and training policy with a mini-batch size of 128.

| method | depth | k | #params | CIFAR-10(%) |
|---|---|---|---|---|
| pre-resnet [10] | 24 | 1 | 0.29M | 7.75 (7.76±0.13) |
| | 68 | 1 | 1.0M | 6.27 (6.33±0.24) |
| | 110 | 1 | 1.7M | 6.02 (6.02±0.11) |
| multi-resnet [ours] | 8 | 4 | 0.29M | 9.28 (9.28±0.07) |
| | 20 | 4 | 1.0M | 6.31 (6.29±0.22) |
| | 30 | 4 | 1.7M | 5.89 (5.85±0.12) |

TABLE II: CIFAR-10 test errors of the multi-residual networks and the original residual networks, where $k$ is the number of functions. The results are in the form of $median$ with $mean \pm std$ in parentheses from five runs.

The Multi-ResNet with 8-layers depth has a test error rate of $9.28\%$, while the original ResNet with 24 layers, and roughly the same number of parameters, has an error rate of $7.75\%$. This is the scenario whereby the network depth is too shallow

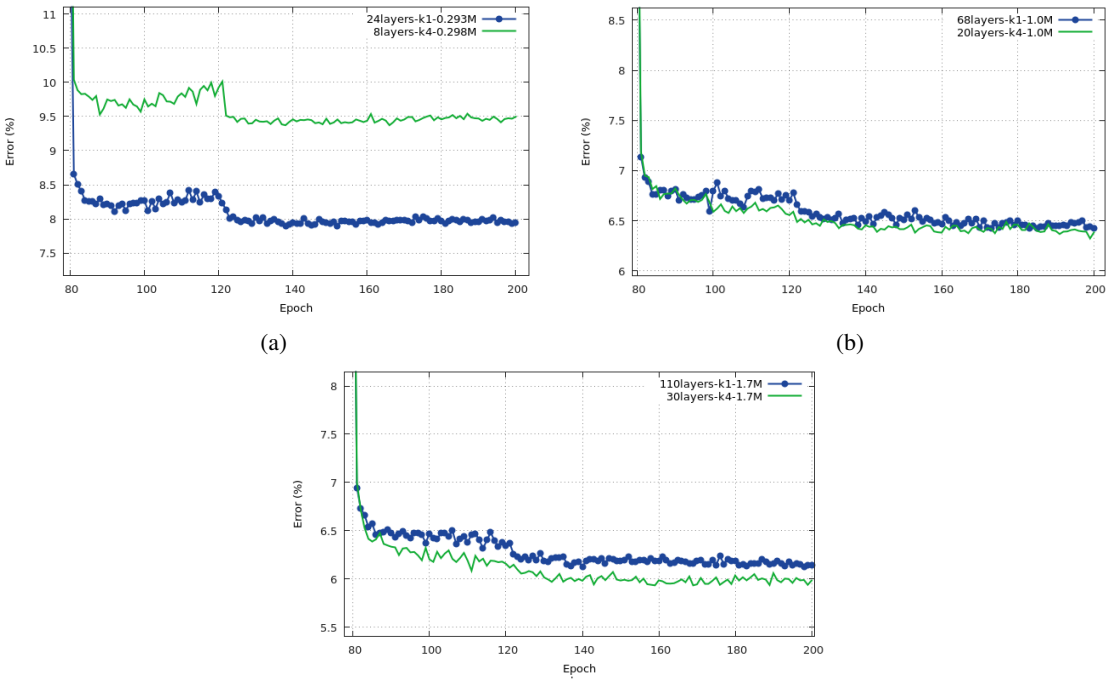


Fig. 3: Comparing residual network and the proposed multi-residual network on CIFAR-10 test set to show the effective range phenomena. Each curve is mean over 5 runs. (a) This is the situation that the network depth $< n_0$ in which the multi-residual network performs worse than the original residual network; (b) Both networks have a comparable performance; (c) The proposed multi-residual network outperforms the original residual network.

$(depth < n_0)$, and the multi-residual network performs worse than the residual network (see Figure 3a). On the contrary, the Multi-ResNet with 20-layers depth achieves $6.31\%$ error rate, which is statistically no different than $6.27\%$ for the 68-layer ResNet. Test curves (Figure 3b) also show that both networks have a comparable performance.

Eventually, a 30-layer deep Multi-ResNet achieves $5.89\%$ error rate. This is slightly better than the 110-layer ResNet that have the error of $6.02\%$ ($6.37\%$ in [10]). Figure 3c also clearly shows that the multi-residual network performance is superior to that of the original residual network. It can be seen that although each pair have almost the same number of parameters and computational complexity, they act very differently. These results support the hypothesis pertaining to the effective range.

In the previous section, we argue that multi-residual network is able to improve classification accuracy of the residual network when the network is deeper than a threshold $n_0$. This effect can be seen in Figure 3. Based on the observations in Table II, for this particular dataset and network/block architecture, the threshold $n_0$ is approximately 20. Furthermore, by increasing the number of functions, better accuracy can be obtained. However, a trade-off has been observed between the network depth and the number of function. This means that, one might need to choose a suitable number of residual functions, and depth to achieve the best performance.

### B. CIFAR Experiments

Table III shows the results of multi-residual networks along with those from the original residual networks and other state-of-the-art models. The networks with $6n + 2$ layers use the basic block with two $3 \times 3$ convolutional layers, and the networks with $9n + 2$ layers use the bottleneck block architecture, which has a single $3 \times 3$ convolutional layer surrounded by two $1 \times 1$ convolutional layers [8]. We also trained wider [36] versions of Multi-ResNet and show that it achieves state-of-the-art performance. One can see that the proposed multi-residual network outperforms almost all of the existing models on CIFAR-10 and CIFAR-100 with the test error rate of $3.73\%$ and $19.45\%$ respectively.

**Complexity of the proposed model.** Increasing the number of residual functions by $k$ increases the number of parameters by a factor of $k$, and the computational complexity of the multi-residual network also increases linearly with the number of residual functions. This results in the memory and computational complexity similar to those of the original residual networks with the same number of convolutional layers [10].

### C. ImageNet Experiments

We also perform experiments on the ImageNet 2012 classification dataset [24]. ImageNet is a dataset containing around 1.28 million training images from 1000 categories of objects that is largely used in computer vision applications. All trainings, in this section, are done using stochastic gradient descent up to 90 epochs. The hyper parameters described

| method | | | | CIFAR-10(%) | CIFAR-100(%) |
|---|---|---|---|---|---|
| NIN[20] | | | | 8.81 | 35.68 |
| DSN[19] | | | | 8.22 | 34.57 |
| FitNet[23] | | | | 8.39 | 35.04 |
| Highway[31] | | | | 7.72 | 32.39 |
| All-CNN[28] | | | | 7.25 | 33.71 |
| ELU[3] | | | | 6.55 | 24.28 |
| method | depth | k,(w) | #parameters | | |
| resnet[8] | 110 | 1 | 1.7M | 6.43$_{(6.61\pm0.16)}$ | 25.16 |
| | 1202 | 1 | 19.4M | 7.93 | 27.82 |
| pre-resnet[10] | 110 | 1 | 1.7M | 6.37 | - |
| | 164 | 1 | 1.7M | 5.46 | 24.33 |
| | 1001 | 1 | 10.2M | 4.62$_{(4.69\pm0.20)}$† | 22.71$_{(22.68\pm0.22)}$ |
| stoch-depth[13] | 110 | 1 | 1.7M | 5.25 | 24.58 |
| | 1001 | 1 | 10.2M | 4.91 | - |
| swapout[27] | 20 | 1,(2) | 1.1M | 6.58 | 25.86 |
| | 32 | 1,(4) | 7.43M | 4.76 | 22.72 |
| wide-resnet[36] | 40 | 1,(4) | 8.7M | 4.97 | 22.89 |
| | 16 | 1,(8) | 11.0M | 4.81 | 22.07 |
| | 28 | 1,(10) | 36.5M | 4.17 | 20.50 |
| DenseNet[12]† | 100 | 1 | 7.0M | 4.10 | 20.20 |
| | 100 | 1 | 27.2M | 3.74 | **19.25** |
| multi-resnet [ours]† | 200 | 5 | 10.2M | **4.35**$_{(4.36\pm0.04)}$ | **20.42**$_{(20.44\pm0.15)}$ |
| | 398 | 5 | 20.4M | 3.92 | 20.59 |
| | 26 | 2,(10) | 72M | 3.96 | **19.45** |
| | 26 | 4,(10) | 145M | **3.73** | 19.60 |

TABLE III: Comparison of test error rates on CIFAR-10 and CIFAR-100. The results in the form of $median(mean \pm std)$ are based on five runs, while others are based on one run. All results are obtained with a mini-batch size of 128, except † with a mini-batch size of 64. The number of residual functions in each residual block is denoted as k, and (w) is the widening factor for wider models.

earlier are used, excluding the learning rate which is divided by 10 every 30 iteration. The networks are trained and tested on $224 \times 224$ crops using scale and aspect ratio augmentation [10], [32].

| method | depth | k | Top-1(%) | 10-Crop(%) |
|---|---|---|---|---|
| pre-resnet | 34 | 1 | 26.73 | 24.77 |
| [10] | 200 | 1 | 21.66 | 20.15 |
| multi-resnet | 18 | 2 | 27.39 | 25.61 |
| [ours] | 101 | 2 | **21.53** | **19.93** |

TABLE IV: Top-1 error rate comparison of deep residual networks and multi-residual networks on ILSVRC 2012 validation set. Multi-residual network outperforms deep residual networks.

Table IV verifies that multi-residual network outperforms a deep residual network with the same number of convolutional layer, as long as the networks are deeper than a threshold. Specifically, the 101-layer Multi-ResNet with two residual functions outperforms the 200-layer ResNet by $0.13\%$ top-1 error rate with the same computational complexity. By testing on multiple crops, the Multi-ResNet outperforms the residual network with $0.22\%$.

Concurrently, ResNeXt [35] and PolyNet [38] obtained second and third place in the ILSVR 2016 classification task with $3.03\%$ and $3.04\%$ top-5 error rate respectively. They are similar to our network architecture in the sense that they both increase the number of functions in the residual blocks. PolyNet also exploits second order paths that compute two functions sequentially in the same block.

### D. Toward Model Parallelism

Although deep residual networks are extremely accurate, their computational complexity is a serious bottleneck to their performance. On the other hand, by simply implementing multi-residual networks, one does not make use of the increase in network width and the reduction in network depth. This is because eventually the residual functions in each residual block are computed and added in a sequential manner. Moreover, the parallel structure of multi-residual networks inspired us to examine the effects of model parallelism as opposed to the more commonly used data parallelism.

Data parallelism splits the data samples among the available GPUs and every GPU computes the same network on its portion of data (Single Instruction Multiple Data), and sends the results back to the main GPU to perform the optimization step. On the contrary, model parallelism splits the model among the desired GPUs and each GPU computes a different part of the model on the same data (Multiple Instruction Single Data) [4], [16]. More precisely, for every multi-residual block with $k$ residual functions, we split the model between two GPUs and each GPU calculates $k/2$ of the residual functions in both forward and backward passes (see Figure 4). The results are then combined on the first GPU to perform the optimization step. Furthermore, the parallelization of each block is believed to reduce the total computational cost of the network.

Using the proposed model parallelism, we compare the computational complexity of the multi-residual network with a similar deep residual network that exploits data parallelism. All experiments are done using Nvidia Tesla K80 GPUs which consists of two sub GPUs connected with a PCI-Express

| pre-resnet [10] | | | multi-resnet [ours] | | | mini-batch | speed up |
|---|---|---|---|---|---|---|---|
| depth | k | Time | depth | k | Time | size | |
| 218 | 1 | 413ms | 110 | 2 | 462ms | 128 | - |
| 434 | 1 | 838ms | 110 | 4 | 804ms | 128 | 4% |
| 650 | 1 | 1284ms | 110 | 6 | 1158ms | 128 | 10% |
| 218 | 1 | 137ms | 110 | 2 | 136ms | 32 | 1% |
| 434 | 1 | 273ms | 110 | 4 | 238ms | 32 | 13% |
| 650 | 1 | 402ms | 110 | 6 | 341ms | 32 | 15% |

TABLE V: Computational time comparison between multi-residual networks and deep residual networks. The captured times are the time for a single SGD step using two GPUs. Multi-residual networks use model parallelism and residual networks exploit data parallelism.

(Gen3) link. This link is capable of transferring data up to 16 GB/s. The elapsed time for a single stochastic gradient descent step including the forward pass, backward pass and parameter update are shown in Table V.

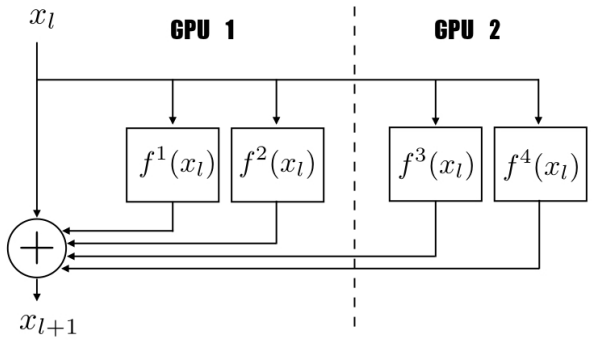


Fig. 4: Model parallelization of a multi-residual block with four residual functions on two GPUs.

In the proposed model parallelism, the inputs and outputs of blocks must be transferred between the GPUs, which occupies most of the computational time. While, in data parallelism, every GPU performs a single forward and backward step independent of others. Nevertheless, Table V demonstrates that the multi-residual network with model parallelism still has less computational complexity than the corresponding residual network. However, this might not be true in some network architectures because of the communication overhead.

Interestingly, this effect amplifies when the number of data samples on each GPU become less than 32, owing to the fact that threads on the current Nvidia GPUs are dispatched in the arrays of 32 threads (called wrap). Therefore, the computational power of GPU is wasted when there are only 16 samples on the GPU. This is sometimes the case in large-scale training, where one has to reduce the batch size in order to fit a larger network in the GPU memory. Also sometimes smaller mini-batch size obtains better accuracy [10].

Consequently, in order to exploit the advantages of both model and data parallelism, one can utilize a hybrid parallelism. As a result, the hybrid parallelism performs data parallelism among four (K80) GPUs and each GPU performs model parallelism internally between the two sub GPUs. This offers up to 15% computational complexity improvement with respect to the deeper residual network.

## VI. Conclusions

Experiments in this article support the hypothesis that deep residual networks behave like ensembles, rather than a single extremely deep network. Based on a series of analyses and observations, multi-residual networks are introduced. Multi-residual networks exploit multiple functions for the residual blocks which leads to networks that are wider, rather than deeper. The proposed multi-residual network is capable of enhancing classification accuracy of the original residual network and almost all of the existing models on ImageNet, CIFAR-10, and CIFAR-100 datasets. Finally, a model parallelism technique has been investigated to reduce the computational cost of multi-residual networks. By splitting the computation of the multi-residual blocks among processors, the network is able to perform the computation faster.

## Acknowledgment

The authors would like to thank the National Computational Infrastructure for providing us with high-performance computational resources. We thank Hamid Abdi and Chee Peng Lim for their useful discussions.

## References

[1] M. Abdi and S. Nahavandi. Multi-residual networks: Improving the speed and accuracy of residual networks. *arXiv preprint arXiv:1609.05672*, 2016.
[2] Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166, 1994.
[3] D.-A. Clevert, T. Unterthiner, and S. Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.
[4] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, A. Senior, P. Tucker, K. Yang, Q. V. Le, et al. Large scale distributed deep networks. In *Advances in neural information processing systems*, pages 1223–1231, 2012.
[5] R. Eldan and O. Shamir. The power of depth for feedforward neural networks. *arXiv preprint arXiv:1512.03965*, 2015.
[6] J. Håstad. Computational limitations of small-depth circuits. 1987.
[7] J. Håstad and M. Goldmann. On the power of small-depth threshold circuits. *Computational Complexity*, 1(2):113–129, 1991.
[8] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.
[9] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1026–1034, 2015.
[10] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. *arXiv preprint arXiv:1603.05027*, 2016.
[11] S. Hochreiter. Untersuchungen zu dynamischen neuronalen netzen. *Diploma, Technische Universität München*, page 91, 1991.
[12] G. Huang, Z. Liu, and K. Q. Weinberger. Densely connected convolutional networks. *arXiv preprint arXiv:1608.06993*, 2016.

[13] G. Huang, Y. Sun, Z. Liu, D. Sedra, and K. Weinberger. Deep networks with stochastic depth. *arXiv preprint arXiv:1603.09382*, 2016.
[14] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
[15] Y. Kim. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*, 2014.
[16] A. Krizhevsky. One weird trick for parallelizing convolutional neural networks. *arXiv preprint arXiv:1404.5997*, 2014.
[17] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
[18] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
[19] C.-Y. Lee, S. Xie, P. Gallagher, Z. Zhang, and Z. Tu. Deeply-supervised nets. In *AISTATS*, volume 2, page 6, 2015.
[20] M. Lin, Q. Chen, and S. Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013.
[21] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014.
[22] R. Mohan. Deep deconvolutional networks for scene parsing. *arXiv preprint arXiv:1411.4101*, 2014.
[23] A. Romero, N. Ballas, S. E. Kahou, A. Chassang, C. Gatta, and Y. Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014.
[24] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
[25] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229*, 2013.
[26] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
[27] S. Singh, D. Hoiem, and D. Forsyth. Swapout: Learning an ensemble of deep architectures. *arXiv preprint arXiv:1605.06465*, 2016.
[28] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*, 2014.
[29] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
[30] R. K. Srivastava, K. Greff, and J. Schmidhuber. Highway networks. *arXiv preprint arXiv:1505.00387*, 2015.
[31] R. K. Srivastava, K. Greff, and J. Schmidhuber. Training very deep networks. In *Advances in neural information processing systems*, pages 2377–2385, 2015.
[32] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015.
[33] A. Veit, M. J. Wilber, and S. Belongie. Residual networks behave like ensembles of relatively shallow networks. In *Advances in Neural Information Processing Systems*, pages 550–558, 2016.
[34] I. Wallach, M. Dzamba, and A. Heifets. Atomnet: A deep convolutional neural network for bioactivity prediction in structure-based drug discovery. *arXiv preprint arXiv:1510.02855*, 2015.
[35] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. *arXiv preprint arXiv:1611.05431*, 2016.
[36] S. Zagoruyko and N. Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
[37] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *European Conference on Computer Vision*, pages 818–833. Springer, 2014.
[38] X. Zhang, Z. Li, C. C. Loy, and D. Lin. Polynet: A pursuit of structural diversity in very deep networks. *arXiv preprint arXiv:1611.05725*, 2016.